The Portfolio Power Theory Revisited: Evidence from Cross-Category Mergers in US Retailing*

Hendrik Döpper
DICE†

October 10, 2023

Abstract

I study 57 cross-category mergers among manufacturers in the US consumer packaged goods retail industry to assess the presence, direction, and size of portfolio effects. In doing so, I exploit differences in the pre-merger bargaining positions of the manufacturers at different retailers. I provide evidence that the manufacturer with the weaker pre-merger bargaining position at a retailer can benefit from increased sales. This increase is driven by changes in quantities, not prices. In addition, I also study the effect on measures of marginal costs and perceived quality. I find that changes in perceived quality help explain these patterns but that marginal costs do not play an important role. Finally, I discuss possible channels that could lead to this result and how these channels are related to the portfolio power theory.

JEL Codes: D22, D43, L11, L4

Keywords: portfolio power theory, portfolio effects, range effects, cross-market mergers, cross-category mergers

---

*I thank Chris Conlon, Alexandra Gibbon, Andreas Lichter, Alexander MacKay, Matthias Mertens, Nathan Miller, Felix Montag, Joel Stiebale, and seminar participants at DICE for their helpful comments and suggestions. Part of this research was conducted during a research stay at the Department of Economics of Harvard University. I thank Elie Tamer for the invitation.
Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.
Computational infrastructure and support were provided by the Centre for Information and Media Technology at Heinrich Heine University Düsseldorf. I gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG) (project 235577387/GRK1974) and the Joachim Herz Foundation.

†Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf. Email: doepper@dice.hhu.de.

1 Introduction

In her speech to the International Competition Network in May 2022, Lina Khan—Chair of the Federal Trade Commission (FTC)—identified three key areas in which merger enforcement in its current form has no bite and where she seeks adjustments in the future, among them the assessment of non-horizontal mergers. While explicitly referring to "deals that might be described as [..] conglomerate", she said that "[w]e must examine how a range of strategies and effects, including extension strategies and portfolio effects, may warrant enforcement action."[1] Her approach to intensifying merger enforcement is part of a larger policy agenda of US President Joe Biden to reverse trends that led to "less competition" and "more concentration" in the previous decades.[2]

Although the portfolio power theory, sometimes referred to as range or portfolio effects, is not new, the literature lacks a clear definition of what it means by these effects.[3] The idea is usually that if two firms sell their products to the same downstream firms, a merger can benefit them even if their product portfolios do not overlap before the merger. In other words, the increase in the sheer size of a firm's product portfolio can change market outcomes, leaving aside possible substitutability and complementarity considerations within the portfolio.

The channel that is often discussed in this context builds on the idea that up- and downstream firms negotiate with each other over terms of supply. These terms of supply can include financial payments, such as linear wholesale prices or lump sum transfers, but can also include non-financial variables. For instance, the downstream firm could spend more effort on promoting and selling the upstream firm's products. A merger can benefit merging upstream firms by a shift in the so-called gains from trade, that is, the additional gain in profit for a bargaining party due to a collaboration with a firm on the other market side. The idea is that a bargaining breakdown becomes increasingly costly for a downstream firm after the merger because the downstream firm now loses access to the products of both upstream firms and not just to those of a single firm. This increases the incentives for the downstream firm to settle the negotiation with the merged upstream entity and gives the merged entity the possibility to demand larger concessions, such as in the form of a larger effort or smaller financial transfers.

[1] A text version of her speech can be found at `https://www.ftc.gov/system/files/ftc_gov/pdf/Remarks%20of%20Chair%20Lina%20M.%20Khan%20at%20the%20ICN%20Conference%20on%20May%206%2C%202022_final.pdf` (last accessed on October 2, 2023).

[2] At the beginning of her speech, Lina Khan herself referred to Biden's agenda, saying, among others: "As you know, competition law in the United States is currently in the midst of a broad and sweeping reassessment. The significance of this reassessment is perhaps best embodied by President Biden's issuance last summer of an Executive Order on Promoting Competition in the American Economy." In his remarks on the executive order, Biden said: "But what we've seen over the past few decades is less competition and more concentration that holds our economy back. We see it in big agriculture, in big tech, in big pharma. The list goes on" (see `https://www.whitehouse.gov/briefing-room/speeches-remarks/2021/07/09/remarks-by-president-biden-at-signing-of-an-executive-order-promoting-competition-in-the-american-economy/`, last accessed on October 2, 2023).

[3] While the OECD (2001) reports a lack of a clear definition, Watson (2003) cites a definition used by the Office of Fair Trading (UK) in a merger case. However, even Watson (2003) acknowledges that "the scope of the concept is somewhat uncertain."

The approach of the current US administration and FTC leadership has reminded some antitrust scholars of the "big is bad" doctrine that played an essential role in pre-Chicago merger enforcement in the 1960s and 1970s and has raised concerns that merger enforcement in the future may be less based on economic theory. One reason for this impression is clearly that Lina Khan herself referred to this period in her speech.[4] While she tried to correct this impression later,[5] part of the perception might also be a strong conflict that raged among antitrust scholars in North America and the EU around the turn of the millennium. At this time, the European Commission used the portfolio power theory in its assessment of multiple merger cases, the most famous one being General Electric/Honeywell. While some US scholars deemed the portfolio power theory to be driven by the pre-Chicago "big is bad" thinking and concluded that the decision by the European Commission to block the merger was not based on economic theory (see, for instance, Evans and Salinger, 2002; Patterson and Shapiro, 2001), others pointed to a wrong understanding of the underlying theory and argued that the analysis of economic models does support the decision (see, for instance, Choi, 2001; Reynolds and Ordover, 2002).

Interestingly, while the General Electric/Honeywell decision led to a heated and (to some extent) also fruitful debate about the potentially pro- and anti-competitive effects stemming from the portfolio power theory, surprisingly little empirical research has been conducted since then to assess the presence, direction, and size of portfolio effects in practice (two notable exceptions discussed below). In light of the latest developments in the US, this paper seeks to revisit the portfolio power theory in the context of the US consumer packaged goods retail industry. I study 57 mergers between manufacturers and analyze the impact on the interactions with the retailers and market outcomes.

The mergers in my sample have two characteristics that make them particularly useful for studying possible portfolio effects. First, they are usually cross-category mergers, meaning that the manufacturers have (almost) no overlap in the product categories in which they were active before the merger. Thus, these mergers would likely be classified as conglomerate mergers by antitrust practitioners and are not affected by horizontal merger effects stemming from a reduction in the number of competitors. Second, many mergers are characterized by strong asymmetries in manufacturers' pre-merger sales to various retailers. I take these pre-merger sales as proxies for the bargaining positions of the manufacturers in negotiations with the retailers. My approach follows the idea that if two manufacturers merge and one manufacturer has a better pre-merger bargaining position with a retailer than the other, the manufacturer

[4] In her last sentence, before referring to the portfolio effects, she said: "While the U.S. antitrust agencies energetically grappled with some of these dynamics during the era of industrial-era conglomerates in the 1960s and 70s, we must update that thinking for the current economy."

[5] The reservations of other scholars resulted not only from her speech but also from a number of other statements and actions on her side. The news article "FTC chair defends track record on antitrust challenges, says big isn't categorically bad" documents one example of her attempt to reverse this impression. It can be found at `https://www.cnbc.com/2023/07/24/ftc-chair-lina-khan-defends-track-record-on-antitrust-challenges.html` (last accessed on October 2, 2023)

with the weaker pre-merger bargaining position may benefit from the merger because the joint bargaining position yields an improvement compared to the pre-merger situation.

I provide evidence that manufacturers with weaker pre-merger bargaining positions tend to benefit from mergers, while manufacturers with stronger pre-merger bargaining positions tend to be harmed. These benefits (and losses) come through increases (decreases) in revenues, which are almost entirely driven by increases (decreases) in the quantities sold and not by changes in prices. To dig further into possible mechanisms behind these results, I then use the work of Döpper et al. (2023) to derive measures for marginal costs and non-price characteristics of the products. I provide evidence that changes in marginal costs do not drive portfolio effects but that changes in the non-price attributes play a crucial role. I link these findings to two possible explanations related to the portfolio power theory: My first explanation builds on the argument outlined above that a merger shifts the gains from trade, which increases the incentives for the retailers to settle negotiations with the merging manufacturers and to make larger concessions to the merged entity. In the consumer packaged goods retail industry, negotiations are usually not just about financial payments but also about the effort that a retailer puts into selling and promoting the manufacturers' products. These efforts can take the form of more or better shelf space or increased in-store promotional activities. If a product is more heavily promoted or better placed on the shelf, this could increase consumers' perception of the quality of the manufacturers' products, leading to a larger number of sales. The second channel is that manufacturers can achieve synergy gains by joining forces in the organization of a joint distribution network. This also increases the incentives for retailers to spend more effort on the products of the merging manufacturers because stockouts (or similar problems) are less likely to occur. Finally, I briefly discuss why two alternative explanations—increased (retailer-independent) advertising spending and efficiency gains beyond the distribution network—are less likely to explain the documented patterns.

For future versions of this paper, I intend to provide a structural model that helps clarify the mechanism behind the documented patterns and assess the implications for welfare and profit sharing among manufacturers and retailers in order to discuss the pro- or anti-competitive nature of the portfolio effects in the context of the US consumer packaged goods retail industry.

Concerning the related literature, two other papers studying the portfolio effects of conglomerate mergers are worth mentioning. Park (2009) and Chunga and Jeon (2014) study four and five mergers between South Korean beer and soju manufacturers, respectively. While the South Korean beer market is dominated by a small number of large manufacturers, past and current regulations have led to a market structure with strong regional players in the soju market (one strong regional player per region). Park (2009) uses a structural demand model to investigate the presence of portfolio effects and finds no evidence for such effects. In contrast, Chunga and Jeon (2014) use a reduced-form approach and a slightly different set of mergers. They provide evidence that large beer manufacturers are able to leverage their size in some regions to push the products of the integrated soju manufacturers to wholesalers if the soju manufacturers did not have a strong position in the region prior to the mergers. The fact that this effect is only present if the soju manufacturers were small competitors in the respective regions before the mergers suggests that the portfolio effect helps to increase local competition and thus may be pro-competitive. This paper differs from the contributions of Park (2009) and Chunga and Jeon (2014) in numerous respects. For instance, the soju market is heavily regulated (for instance, ban on wholesale price discrimination and ban on TV and radio advertising), while most of the product categories in my study experience rather little regulation (if any). In addition, I consider a much broader set of mergers as well as a large number of product categories, and while the aforementioned studies only analyze the impact on market shares, I consider various other market outcomes.

My study also contributes to the literature on cross-market mergers, that is, mergers between firms that operate in different (geographic or product) markets and, therefore, would usually not raise concerns by antitrust authorities. Cross-market mergers have recently attracted the attention of scholars in health economics. Lewis and Pflum (2017) and Dafny et al. (2019) provide empirical evidence that cross-market hospital mergers can impact market outcomes and lead to price increases. Both studies are similar to this paper in that they provide an in-depth analysis of the interplay between merging upstream firms and downstream intermediaries that bundle the upstream products. However, this paper differs from these studies in that it focuses on a lack of overlap in product rather than geographic markets, deals with a different industry, and documents the effects of cross-market mergers on revenues that are driven by changes in quantities rather than prices (and thus by changes in the non-price characteristics of the products).

Another strand of literature that has similarities to the one on cross-market mergers deals with cross-border mergers, that is, mergers of firms located in different countries. While some of these mergers may also be affected by a market overlap, others are not, so the literature on cross-border mergers is often concerned with discussing merger effects that arise in the absence of overlapping (geographical) markets. For instance, Guadalupe et al. (2012) study the impact of cross-border acquisitions on Spanish manufacturing firms and find that acquired firms' innovation activities increase post-merger. One channel that they identify is that the acquired firms gain better access to foreign markets through their new parents. This is similar to one of the channels that I discuss in Section 5, where the acquired targets benefit in negotiations with the retailers and are able to increase the effort provided by the retailers.

Finally, this paper also contributes to recent discussions about the effectiveness of antitrust enforcement in the EU and the US. Bhattacharya et al. (2023) use the same scanner data as in this study (NielsenIQ) in combination with the SDC Platinum merger database from Thompson Reuters to analyze the effects of mergers that might be potentially relevant for antitrust authorities.[6] They use a structural model that allows them to evaluate counterfactual scenarios in which they can vary the intensity of merger enforcement and find that an increase in the intensity would lead to a substantial reduction in Type II errors, however, at the expense of a much larger number of cases to be examined. In another study, Affeldt et al., 2021 focus on potential efficiency gains from mergers, which are often used as a defense against potential merger remedies and prohibitions. They conclude that "[c]ompensating efficiencies appear to be simply too large to be achieved by real world mergers [..]." My paper fits into this strand of literature in that I provide evidence for the existence of merger effects that are often ignored by antitrust authorities. If these effects benefit consumers, they could be used by the merging firms as an additional defense tool. If, in contrast, these effects harm consumers, antitrust authorities might want to block an even larger number of mergers.

The remainder of this paper is structured as follows: Section 2 describes the data. Section 3 introduces three important definitions (Subsection 3.1), provides insights about the cross-category activities of the manufacturers (Subsection 3.2), and documents the effects of cross-category mergers on directly observable outcomes like revenues, quantities, and prices (Subsection 3.3). Since the scope of directly observable market outcomes is limited, I then use the work of Döpper et al. (2023) to shed some light on other measures like marginal costs in Section 4. In doing so, I first describe the model (Subsection 4.1) and the empirical strategy (Subsection 4.2) before extending my analysis of the effects of cross-category mergers in Subsection 4.3. Finally, I discuss possible mechanisms that can drive these results in Section 5 before summarizing my main findings in Section 6.

2 Data

A common problem of empirical studies of vertical chains is that contracts between up- and downstream firms are typically not observed, and data on the vertical relations is missing. Therefore, most of the IO literature combines structural models based on assumptions about firm conduct with data on consumer behavior. My analysis follows this approach and uses two widely used data sets for the US consumer packaged goods retail industry that are provided by NielsenIQ in collaboration with the Kilts Center for Marketing Data Center at the University of Chicago. Both data sets provide information about the consumers' purchasing decisions in a large variety of product categories.[7] The product categories cover both food and non-food products, such as ready-to-eat cereals, shampoo, and bottled water. The difference between the data sets stems from the source from which the data originates.

[6] Another study that uses the NielsenIQ data set in combination with SDC Platinum is Majerovitz and Yu (2023). The authors focus on the average horizontal merger, which is characterized by strong asymmetries, typically including a small target and a large acquirer.

[7] NielsenIQ distinguishes between three different product group classifications. I follow Döpper et al. (2023) and use the so-called product modules as an approximation for the product markets. I will refer to these product markets as (product) categories.

The first data set—the so-called Retailer Panel—is directly reported by a large set of US retailers. Each retailer provides weekly sales information for its stores. The sales are reported at the level of bar codes where a bar code is defined by the Universal Product Code (UPC). The sales information is complemented with additional information about store, retailer, and product characteristics. The retailers can be categorized into different retail channels. For this analysis, I restrict my attention to food stores, mass merchandisers, and drug stores[8].

The second data set is the so-called Consumer Panel and contains information about shopping trips of individual households. The households in the sample participate in a program operated by NielsenIQ and report the data themselves. The information about the purchase of a product is complemented by additional information about household, shopping trip, and product characteristics. The households can be reweighted so that they are representative of the US population with respect to a number of observable demographic characteristics.

The different data sources come with different advantages and disadvantages. The Retailer Panel is a useful starting point for my analysis since it covers a large portion of the total household spending in the industry. It does, however, not contain information about the relationship between household characteristics and purchasing decisions because sales are aggregated at the store level. Therefore, it seems reasonable to complement the Retailer Panel with the Consumer Panel if information about individual factors is required.

Döpper et al. (2023) use this strategy to analyze the evolution of market power in the US consumer packaged goods retail industry. To this end, they estimate BLP-style demand systems for 133 product categories between 2006 and 2019. I follow their approach in the sense that I perform the same steps to process the raw NielsenIQ data sets and adopt their estimation strategy to gain insights beyond what can be learned from directly observable measures. More precisely, their approach allows me to recover a measure for marginal costs and a metric to quantify the impact of product characteristics other than the price on consumers' decision-making.

Döpper et al. (2023) use the Retailer Panel to calculate product-level market shares across different regions[9] and retail outlets. Their analysis focuses on 133 product categories, and in each product category, they aggregate the data along three dimensions. First, they choose a different product definition than the UPCs and aggregate sales to the brand level. The reason is that UPCs are often very narrowly defined and do not correspond to what the consumer perceives as a product.[10] For instance, there can be different package sizes of a product and each package size can have a different UPC. This leads to a large number of UPCs, which

[8]These are the retail channels for which NielsenIQ provides good coverage for all years. The other product channels that I exclude are dollar, club, convenience, and liquor stores.

[9]Döpper et al. (2023) focus on the 22 largest Designated Market Areas, which are coherent areas defined by Nielsen based on media markets. The idea is that consumers in each region are exposed to the same marketing campaigns because they are served by (almost) the same newspapers and TV and radio stations.

[10]Döpper et al. (2023) provide a list of examples of what brand names look like in the dataset (for instance, see their footnote 12).

makes it difficult to infer cross-substitution patterns in practice. Aggregating to the brand level substantially lowers the number of products in a category and allows to circumvent problems related to the large number of alternatives. Döpper et al. (2023) further restrict their attention to the 20 top-selling brands and consolidate the remaining brands into a fringe brand. Second, they aggregate the sales across multiple stores of a retailer in a region. This allows to reduce the likelihood of zero market shares (as discussed in, among others, Dubé et al., 2021 and Gandhi et al., 2023) and to keep the data set at a manageable size. Third, they consolidate weekly sales into quarterly sales. This also serves the purpose of a lower likelihood of zero market shares and, in addition, allows to better account for potential concerns arising from the stockpiling behavior of households (as discussed in, among others, Hendel and Nevo, 2006). The three aggregation steps lead to a data set where, for each product category and year, one observation is provided for each brand sold at a retailer in a region in a quarter. For my reduced-form regressions, I further aggregate the data across regions and quarters so that I have brand-retailer-specific metrics for each category and year.

This consolidated data set lacks a link between household characteristics and consumers' purchasing decisions. However, this link is important to account for heterogeneity in consumers' responses to price differences and changes, thereby preventing the estimation of rich substitution patterns. Döpper et al. (2023) use the Consumer Panel in two ways to add this heterogeneity component to the data. First, they calculate the annual distribution of household characteristics at the regional level. In doing so, they restrict their attention to two characteristics, namely the household income and a variable indicating whether a household has children or not. Second, they calculate so-called micro-moments that are used to capture heterogeneity in the target audience of the brands. A micro-moment corresponds to the average characteristic of a consumer buying a certain brand. Döpper et al. (2023) calculate micro-moments for all brands in all product categories and allow them to vary across regions and time.

Finally, three other data sets complement the NielsenIQ data. First, Capital IQ provides a snapshot of ownership information that allows to link brands to manufacturers. Based on this, the Zephyr merger database allows to identify mergers in the sample and keep track of changes in ownership over time.[11] Finally, Döpper et al. (2023) use a Consumer Price Index (CPI) to deflate all monetary measures (like prices). The CPI[12] used in the analysis excludes most of the product categories in the sample so that changes in monetary measures can be interpreted (roughly) as relative to changes in the prices of other goods in the economy.

[11]The compilation of ownership information in Döpper et al. (2023) is not ideal for my analysis because I do not have information about the owners of the brands that are collapsed into the fringe brand. In addition, information about changes in ownership is available only at the annual level but not at the quarterly level. I am currently working on more detailed ownership information so that this problem is likely to be fixed in future versions of this paper.

[12]The CPI used in the analysis is the "Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average". See `https://fred.stlouisfed.org/series/CPILFESL` (last accessed on October 2, 2023) for details.

3 Cross-Category Activities and Mergers

3.1 Definitions

Cross-category activities of firms are at the core of my analysis. To avoid any confusion about what I mean by cross-category activities or cross-category mergers, I introduce three definitions. I start with a terminology that describes the activities of firms in two or more product categories.

Definition 1. *A firm is said to be "active cross-category" if its products belong to more than one product category.*

The primary objective of this definition is to describe the activities of manufacturers. The reason is that—as I will demonstrate below—there is large heterogeneity in the firms' product assortment on the manufacturers' side. The definition is, however, also applicable to retailers. All retailers are active in a large variety of product categories and can thus be considered as being active cross-category.

Next, I turn the focus to mergers.

Definition 2. *If two or more firms merge, the merger is said to be a "cross-category merger" if at least one product category exists in which only one of the merging parties was active prior to the merger.*

My analysis solely focuses on mergers of manufacturers. Based on the definition, I can attach a label to each merger indicating whether it is cross-category. This may, however, not be informative about how a merger affects a single product category. Think of two firms, with firm 1 being active in the categories A and B and firm 2 being active in the categories A and C. Both firms are active in more than one product category, therefore they are active cross-category according to Definition 1. In addition, a merger between these two firms would be called a cross-category merger. The reason is that there is only one merging firm active in each of the categories B and C prior to the merger. Although the merger would be a cross-category merger in my terminology, the merger may also generate effects by reducing competition in some product categories where the assortment of the merging parties overlapped before the merger. In my example, this would be category A. To better describe the impact of a merger on a particular product category, I introduce the following definition.

Definition 3. *A product category is affected by a merger if at least one of the merging parties was active in the category prior to the merger. If only one merging party was active in the category before the merger, I say that the category was affected "cross-category." Otherwise, I say that the category is affected "horizontally."*

The term "horizontally" refers to the terminology of a "horizontal merger" and is frequently used in the literature to describe a merger between two or more firms in the same market, which leads to a reduction in the number of competitors and, thus, usually also in competition. As I will discuss later, the set of cross-category mergers used in my analysis contains some cross-category mergers that also affect categories *horizontally*, but the number of categories is rather small. Therefore, I will simply exclude these merger-category combinations in my analysis and focus on the remaining categories without horizontal effects.

3.2 Cross-Category Activities

Table 1: Cross-Category Activities

Panel A: Manufacturers

Number of Firms		Number of Categories per Firm							Sales Share of Firms
Total	Cross-Category	Mean	Median	75% P.	90% P.	Largest 3			Active in 4+ Categories
743	223	2.14	1	2	4	29	31	40	44.15

Panel B: Retailers

Number of Firms		Number of Categories per Firm				
Total	Cross-Category	Mean	10% P.	25% P.	Median	75% P.
101	101	126.68	125	129	132	133

Although the focus of my analysis is on cross-category mergers, it seems reasonable to establish some facts about cross-category activities first. Panel A of Table 1 provides some basic statistics about the activities of the manufacturers. It shows that out of the 743 firms in my sample, around 70% are active in only one product category. In other words, the average (median) firm is not active cross-category.

This is also visible from the distribution of the number of categories per firm, which shows a median of one. The number increases only slightly to 2 and 4 at the 75^{th} and 90^{th} percentile, respectively. Given that I have 133 categories in my sample, these numbers can be considered small. This highlights that the remaining 223 cross-category firms are typically active in a small number of categories.

The large majority of firms that are active in only a few categories is accompanied by a small set of large firms. These firms can be of substantial size. For example, the product assortment of the three largest firms spans 29, 31, and 40 categories. Although these numbers are very large compared to the percentiles listed in Panel A of Table 1, it is important to keep in mind that they represent only 22%, 23%, and 30% of the universe of categories in my sample.

Contrary to the manufacturers and as visible from Panel B of Table 1, the retailers' assortments typically cover a large portion of the categories in my sample. As mentioned earlier, all 101 retailers are active cross-category. The median retailer covers all categories except one, and even the 10^{th} percentile of the distribution of the number of categories per retailer is 125, which represents almost 94% of the categories in my sample.

From an economic perspective, the two panels of Table 1 stress the importance of examining to what extent cross-category effects (and thus portfolio effects) play a role in bargaining and how they shape the relationship between manufacturers and retailers. For instance, if cross-category effects are absent, bargaining outcomes solely depend on the market positions of the firms in a given category (like their market size or their brand valuations). In other words, if a manufacturer is active in a single category and holds a strong market position, it will also have a strong bargaining leverage over the retailers. If, in contrast, cross-category effects are extremely important, the bargaining leverage of such a manufacturer can be expected to be almost negligible. Even the biggest manufacturer in my sample would have a rather weak bargaining position because it is active in "only" around 30% of the categories.

Panel A of Table 1 is useful to get a first impression of the cross-category activities of the manufacturers. It is, however, not per se informative about how these activities look like within categories. It could, for instance, be the case that most of the cross-category firms cluster in a small number of categories while other categories are almost unaffected by cross-category firms. The purpose of Table 2 is to show that this is indeed not the case and that cross-category activities are a widespread phenomenon.

Table 2: Cross-Category Activities of Manufacturers by Product Category

		Number of Firms per Year		Share of Revenues		
Rank	Product Category	Total	Cross-Category	Top 20 Brands	Cross-Category Firms	Private Labels
1	Cereal - Ready to Eat	6	4	0.56	0.48	0.08
2	Candy - Chocolate	7	4	0.52	0.42	0.03
3	Candy - Non-Chocolate	12	5	0.57	0.35	0.09
4	Deodorants - Personal	8	6	0.79	0.78	0.00
5	Soap - Specialty	10	6	0.69	0.61	0.05
6	Tooth Cleaners	5	4	0.74	0.74	0.00
7	Shampoo - Liquid/Powder	9	5	0.60	0.53	0.03
8	Cookies	8	5	0.63	0.46	0.16
9	Sanitary Napkins	5	3	0.75	0.62	0.13
10	Cold Remedies - Adult	10	6	0.88	0.45	0.28
20	Bottled Water	10	7	0.88	0.65	0.22
40	Baby Formula	5	2	0.80	0.37	0.04
60	Nuts - Bags	17	10	0.86	0.42	0.32
80	Fresh Muffins	13	7	0.92	0.71	0.19
100	Tuna - Shelf Stable	14	7	0.99	0.85	0.11
120	Cream - Refrigerated	13	10	0.92	0.46	0.45
130	Frozen Poultry	15	6	0.93	0.34	0.51
133	Fresh Mushrooms	17	2	0.96	0.02	0.44
	Mean Values	12	6	0.85	0.59	0.16

Table 2 presents information about the cross-category activities for a subset of categories. The selection of categories is taken from Table 1 in Döpper et al. (2023), and categories are sorted by the number of observations. The value in the first column is the rank resulting from this sorting exercise. The first group of categories (up to the horizontal rule) contains the ten largest categories, while the second part includes a subset of the remaining categories. The last row shows statistics for the average category.

I first focus on the number of firms that are active in the category (column 3) and the corresponding number of cross-category firms (column 4). As indicated in the last row, half of the firms are active cross-category in the average category. Across categories, this ratio varies substantially, but the number of cross-category firms is usually well above zero. Notable exceptions exist in the categories “Baby Formula” and “Fresh mushrooms,” where only two cross-category firms are active. In the first case, this is not surprising given that only five firms are active in total, while in the latter case, the category seems indeed to be less affected by cross-category activities.

The number of firms gives a first impression about the activities of cross-category firms within categories, but it may hide important information because this measure treats all firms equally. An alternative would be to look at the share of revenues that is captured by the brands of cross-category firms (column 6). By construction, these brands are a subset of the leading 20 brands in each category; hence, I also report the revenue share of these 20 brands as a benchmark (column 5). The table shows that the 20 brands account for 85 percent of the revenues in the average category. The subset of brands owned by cross-category firms accounts for about 70 percent of this share and almost 60 percent of the total revenues. This means that cross-category firms tend to be large not only because they offer products in multiple categories but also because their market coverage within a given category is large.

It is worth noting that the revenue share of the leading 20 brands may also include the sales of private labels (see column 7 for the corresponding revenue share). Although private label products are typically treated as being produced by the retailers in IO models, and retailers are cross-category firms, I do not treat private label products as products sold by cross-category firms in my analysis. If I included the 16 percent that private labels account for in the average category, around 88 percent of the revenue share of the leading 20 brands would be associated with cross-category firms.

Finally, Table 2 shows that across all categories, the revenue share of brands sold by cross-category firms almost never drops below one-third and is often substantially larger. A notable exception is the category “Fresh Mushrooms” where cross-category firms account for only 2 percent of the revenues. This is consistent with the initial inspection based on the number of firms.

3.3 Cross-Category Mergers

The previous subsection shows that cross-category activities of both manufacturers and retailers are a widespread phenomenon in the US consumer packaged goods retail industry. In this section, I will explore the effects of cross-category mergers on directly observable market outcomes such as revenues, quantities, and prices. I start my exploration by looking at some statistics that describe the mergers in my data.

Panel A of Table 3 shows that out of the 139 mergers in my sample, 95 can be classified as cross-category mergers. Among these cross-category mergers, 57 mergers are suitable for my analysis. I will refer to them as the baseline sample. The difference between the total number of cross-category mergers and the baseline sample is mostly driven by missing data on the acquirer side. In 32 cases, the acquirers are not active in any category in my sample. These mergers still constitute some form of cross-category mergers since the acquirers are not active in the same product categories as the targets, and hence, these mergers do not reduce competition in these categories. However, since my analysis requires information about both merging parties, I restrict my attention to the baseline sample.[13]

Table 3: Overview of Cross-Category Mergers

Panel A: Number of Mergers

All Mergers	Cross-Category Mergers	
Total	Total	Baseline
139	95	57

Panel B: Characteristics of Cross-Category Mergers

	All Mergers	Per Merger			
	Unique Values	Mean	25% Q.	50% Q.	75% Q.
Targets	55	-	-	-	-
Acquirers	36	-	-	-	-
Categories (Cross-Category)	115	9.56	3	8	13
Categories (Cross-Category, Target)	68	1.88	1	1	2
Categories (Cross-Category, Acquirer)	105	7.68	2	5	12
Categories (Horizontal)	27	0.47	0	0	1
Brands (Target)	140	2.74	1	2	3
Brands (Acquirer)	643	19.91	5	10	22
Total Sales (Target)	-	34.13	3.24	12.35	31.61
Total Sales (Acquirer)	-	325.66	36.84	97.60	423.22
Avg. Sales Share in Category (Target)	-	0.05	0.01	0.02	0.05
Avg. Sales Share in Category (Acquirer)	-	0.09	0.03	0.06	0.12

Panel B of Table 3 provides an overview of the characteristics of the mergers. The second

[13]There might also be different reasons why firms merge. If an acquirer is not active in any of the 133 categories in my sample, I cannot be sure that this firm is active in the consumer packaged goods retail industry at all. For instance, the acquirer could also be a private equity firm.

column lists the number of unique values for a variety of characteristics across all baseline mergers. The remaining columns describe the distribution of these characteristics across mergers. Starting with column 2, it shows that the 57 cross-category mergers involve 55 unique targets. This means that almost all targets were bought only once during the 14 years of my sample period. In contrast, the set of acquirers is smaller and consists of only 36 firms, showing that some manufacturers acquire multiple targets. In fact, there are 14 acquirers that conduct more than one acquisition, and the most active acquirer buys four targets over the 14 years of my sample period.

Cross-category mergers can affect market outcomes in categories on both the target and the acquirer side. In total, 115 categories are affected by at least one merger on either side, with 68 categories being affected at least once on the target side and 105 categories at least once on the acquirer side. This suggests that the acquirers are active in more categories than the targets. This is also visible from the distribution of the number of categories per merger (rows 3 to 6). While the average merger affects about 9.5 categories, only about two categories are affected on the target side, and the remaining approximately 7.5 categories are affected on the acquirer side. This pattern also holds true for the three other percentiles reported in the table.

Row 6 shows that some mergers also affect categories horizontally, that is, both firms are active in these categories before the merger. However, the number of categories is rather small. While there are 27 categories being affected horizontally in total, the average (median) merger shows no overlap in product markets, and even the 75^{th} percentile is only 1.

The fact that acquirers operate in many more categories than the targets could also mean that the targets are much smaller than the acquirers. There is, however, a caveat to this idea: the activity of a firm in a category is, per se, not informative about its success within this category. It could, for instance, be the case that an acquirer is active in many more categories but that its brands target niche consumer segments and realize only small market shares, while the target is highly specialized in a single category but is able to capture a large market share. The remaining rows of the table are used to reject this alternative explanation and to show that the acquirers are indeed much larger than the targets. The rows present three different measures to better capture the full extent of what can be described as "being larger": the total number of brands, the total revenues, and the average revenue share in a category. All three measures point to the fact that acquirers are larger. For instance, the average acquirer has more brands (about 20 vs. 3), realizes larger revenues (about 325 vs. 35 million USD), and captures a larger revenue share within a category (about 9 vs. 5 percent). This pattern does not only hold for the average merger but remains valid when looking at an alternative measure for the average (median instead of mean) and different percentiles of the distribution (25^{th} and 75^{th}).

The asymmetry between targets and acquirers provides further guidance for how I can carry out the analysis of merger effects. Recall the idea of the portfolio power theory that a cross-category merger can benefit the merging parties through an improvement in their bargaining

Figure 1: Histogram of the Pre-Merger Sales Ratio

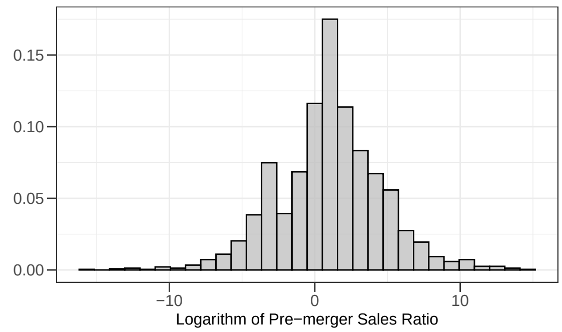

position. If the merging firms are highly asymmetric, the shift in the bargaining position is likely larger for the smaller firm because this firm generated only small revenues before the merger and thus was highly dispensable for the retailers. In contrast, the larger merging party generated large revenues already prior to the merger, and its size increased only marginally through the merger. Therefore, the importance of its assortment does not change a lot from a retailer's perspective. In conclusion, this means that my analysis should be primarily concerned with the effects of mergers on the outcomes of the smaller firms, that is the targets. In addition, I will use the fact that firms' activities vary substantially across retailers.

Consider a brand j belonging to a target. I use $f_j(\tau)$ to denote the ownership of brand j at point τ. The time variable τ is measured in event time; that is, it takes the value 0 in the year of the merger. Thus, $f_j(-1)$ and $f_j(0)$ refer to the independent target before the merger and the acquirer after the merger. The key metric of my analysis is the ratio of the revenues of the acquirer relative to those of the target in the year before the merger.

$$\log\left(\frac{total\ sales_{f_j(0),c,-1}}{total\ sales_{f_j(-1),c,-1}}\right) \tag{1}$$

The indices $f_j(\tau)$ and $f_j(0)$ refer to target and acquirer, while the additional index -1 refers to the last pre-merger period (in event time). The remaining index c denotes the retail chain at which the revenues are generated. Note that the index j belongs only to the ownership variable f_j, but does not enter the total revenues as an additional subscript. This means that the revenues refer to the total revenues of the corresponding firm at retailer c and time -1 and not just those of brand j.

Figure 1 shows a histogram of the pre-merger revenue ratios (in logarithm). I keep the observations at the target-acquirer-retailer level, leaving aside that my analysis will take place at the brand level. The figure shows that most of the observations have a positive value. In fact, the 33^{rd} percentile is about -0.01, indicating that about two-thirds of the observations are positive. The mean (1.03) and median (1.07) are both very similar and close to 1, supporting the fact that the distribution looks rather symmetric.

One important observation from Figure 1 is that while most of the observations are positive, there is still a substantial fraction that is negative (about one-third). This is one of the two reasons that motivates the use of the logarithm in Expression (1) and the subsequent analysis. If the logarithm is negative, the ratio of the revenues must be smaller than one, meaning that the target's revenues at retailer c exceed those of the acquirer. The use of the logarithm allows for opposing effects; that is, the effect is negative if the target's revenues are larger and positive if the target's revenues are smaller. While I impose this relationship by assumption, it is supported by my analysis later (see details below). Another advantage of the logarithm is that it alters the interpretation of the regression coefficients in a meaningful way, allowing me to consider percentage changes in the ratio rather than level changes.

With measure (1) in hand, I can now state the main specification.

$$X_{jc\tau} = \alpha_{jc} + \gamma_{year(\tau)} + \sum_{\ell \in [\underline{\tau}, \overline{\tau}]} \left[\beta_{1\ell} + \beta_{2\ell} \cdot \log \left(\frac{total\ sales_{f_j(-1),c,-1}}{total\ sales_{f_j(0),c,-1}} \right) \right] \cdot D\left(\tau = \ell\right) + \varepsilon_{jc\tau} \quad (2)$$

The variable X_{jct} will be the outcome of interest. For now, this will be the logarithm of either the revenues, quantities, or prices. The indices show the level of aggregation. Observations are at the brand-retailer-event time level, which means that, as noted in Section 2, I abstract from regional differences and only exploit the variation across retailers.

As it is common practice in the event study literature, the specification follows a two-way fixed effects design where α_{jc} captures the brand-retailer fixed effects and $\gamma_{year(t)}$ the year fixed effects. The sum operator loops over all event time periods in a time window from 5 years before to 5 years after the merger. Observations before and after this time window are collapsed into two additional bin categories ($\tau < -5$ and $\tau > 5$).[14] The variable $D(\cdot)$ is a dummy variable taking value 1 if the condition in brackets is satisfied and 0 if not. This means the variable takes value 1 if the observation belongs to the event time period ℓ. The coefficients $\beta_{1\ell}$ and $\beta_{2\ell}$ are supposed to capture the effects of the merger. My sample will include only the brands of the targets. Hence, the idea is to compare treated to not-yet-treated brands, assuming that after accounting for the fixed effects and in absence of the treatment, their developments would follow similar trends. The coefficients $\beta_{1\ell}$ capture effects that are common to all brand-retailer combinations, while the coefficients $\beta_{2\ell}$ capture additional effects arising from pre-merger differences in the bargaining positions of the target and the acquirer at retailer c. In practice, the $\beta_{1\ell}$ coefficients turn out to be usually insignificant and close to zero, so I will treat them as an additional set of fixed effects.[15] The $\beta_{2\ell}$ coefficients in the pre-merger periods are supposed to be 0, with the coefficient in the last pre-merger period ($\ell = -1$) being

[14] I omit the estimates for the bin categories in the following figures. The estimates usually fit the patterns shown in this paper. For my main Figure 2, I also report the estimates in the table in Appendix A.

[15] Appendix D compares the results of Figure 2 (black) to a version where I omit the $\beta_{1\ell}$ coefficients in the specification (gray). The results are similar.

Figure 2: Changes in Sales, Quantities and Prices

(a) Sales

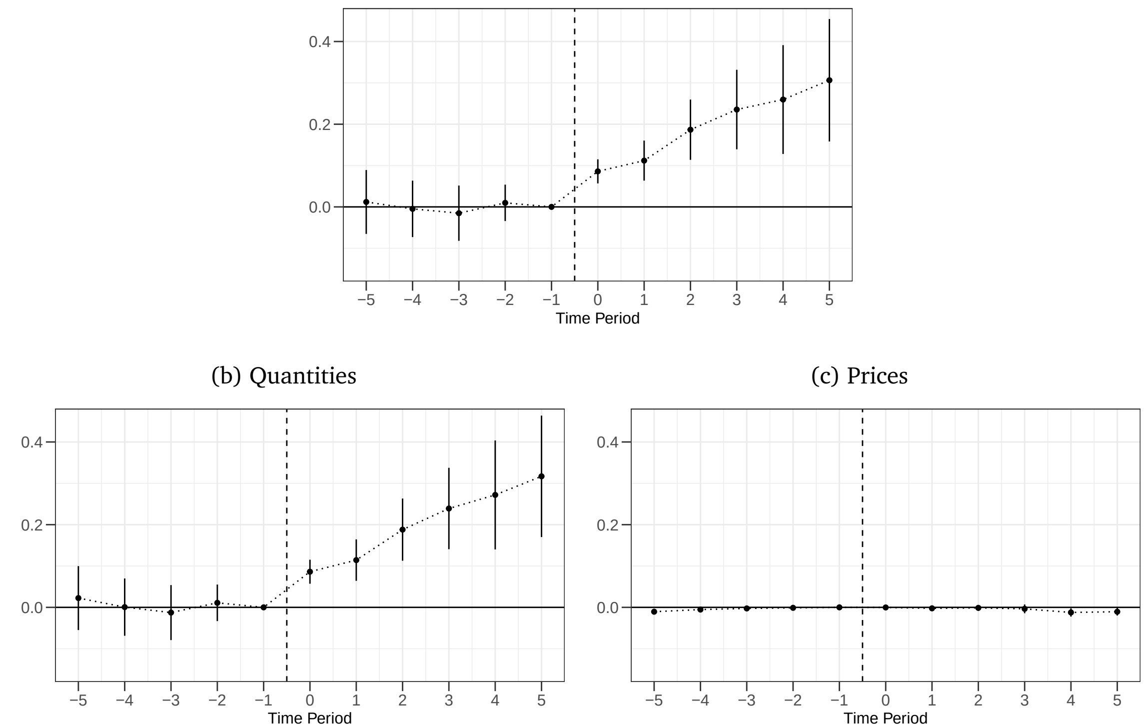

(b) Quantities

(c) Prices

normalized to 0.

The event study literature is currently undergoing significant developments. With this in mind, I will interpret the estimated $\beta_{2\ell}$ coefficients as correlations for now, which may give a first impression of possible cross-category effects. I will discuss a more sophisticated approach based on the recent event study literature and the underlying assumptions later. However, the results will be broadly consistent with the patterns documented based on the initial inspection of the correlations.

Figure 2 visualizes the results when I estimate the main specification with Ordinary Least Squares.[16] It shows the $\beta_{2\ell}$ coefficients for the different time periods. The vertical bars show the 95% confidence intervals, with standard errors being clustered at the merger level. The three panels refer to the three different measures of interest. It is clearly visible that the estimated coefficients in the pre-merger time periods are close to zero, independently of the measure under consideration. This provides some support for the idea that treated and not-yet-treated brands undergo similar developments before the merger and after accounting for the fixed effects.

The estimated coefficients for the merger periods are similar for revenues (Panel (a)) and

[16]Appendix A provides the corresponding table with the estimates.

quantities (Panel (b)). This refers to both the direction and the magnitude of the estimated coefficients.[17] In contrast, the coefficients of the prices are very close to zero. Although some of the coefficients are, in fact, statistically significantly different from 0, they can be considered economically negligible.

The post-merger coefficients for revenues and quantities are positive and increasing over time. Since both the dependent variable and the pre-merger revenue ratio are in logarithms, the interpretation of the coefficients relates to changes in percent. If the pre-merger revenue ratio increases by 1%, ceteris paribus, the revenues (and quantities) of the target's brands at the corresponding retailer increase by about 0.1% in the first year after the merger. This effect increases to about 0.3% 5 years after the merger.

To get a sense of the total effect size, consider the 25^{th}, 50^{th}, and 75^{th} percentile of the distribution of the logarithms of the pre-merger revenue ratios as depicted in Figure 1, which are -0.92, 1.07, and 3.07, respectively. The benchmark is a target that has the same pre-merger revenues at a retailer as the acquirer so that the logarithm of the pre-merger revenues ratio is 0 and there are no cross-category effects. Compared to this benchmark, the cross-category effects lead to a change in revenues (quantities) by roughly -25%, 39%, and 156% (-25%, 40%, and 165%), ceteris paribus, when evaluated at the three percentiles and at the point estimate.[18] These values indicate that the cross-category effects can reach considerable magnitudes.

So far, and as noted above, the analysis documents correlations. The next step is to provide additional evidence that the effects are also causal. To this end, I make use of recent developments in the event study literature. I start my investigation of a potentially causal relationship by discussing the assumptions that would be required for the above analysis to reveal a causal relationship.[19] With the knowledge of which assumptions are unlikely to hold, I can then look for an alternative approach.

I begin with three assumptions that I deem to be unproblematic. The first assumption is the "stable unit treatment value assumption" (SUTVA), which says that the treatment status of one firm does not impact the market outcomes of another firm's brands. In particular, there are no spillovers across firms in my sample. This assumption may seem critical at first glance because mergers clearly impact all firms in the markets in which the merging firms are active.

[17]Appendix B shows the result of the exercise when I use revenue and quantity shares of the brands within the product categories as a dependent variable. The results are similar. The result also remains intact if I consider only a balanced panel with a three-year time window. This is visible from Appendix C that compares the outcomes of Figure 2 (black) with those of a balanced panel (gray).

[18]The formula to calculate the effect size is $100 \cdot \left(\exp\left(\widehat{\beta}_{25} \cdot \text{ratio}_1\right) - 1\right)$, where $\widehat{\beta}_{25}$ is the point estimate, and ratio_1 is the logarithm of the pre-merger revenue ratio. To derive this formula, let $i = 1$ refer to the case where the logarithm of the pre-merger revenue ratio is given by one of the percentiles listed in the text, and $i = 0$ denotes the benchmark case where the logarithm of the ratio is 0. Let ratio_i denote the logarithm of the pre-merger revenue ratio and y_i the variable of interest. The formula results from the following consideration $\log(y_1/y_0) = \log(y_1) - \log(y_0) = \beta_{25} \cdot \text{ratio}_1 - \beta_{25} \cdot \text{ratio}_0 = \beta_{25} \cdot \text{ratio}_1$ and thus $y_1/y_0 - 1 = \exp(\beta_{25} \cdot \text{ratio}_1) - 1$.

[19]The following discussion draws primarily on the survey of Roth et al. (2023). However, due to the fast progress in the event study literature and the importance of this literature for many fields in economics, there are many other good surveys available. Another notable one is de Chaisemartin and D'Haultfœuille (2023).

However, my analysis compares treated to not-yet-treated firms. In particular, I do not follow other papers studying mergers (like, for instance, Ashenfelter and Hosken, 2010) and do not use competing brands or private labels for the comparison. This means that the firms in my sample are usually active in different product categories, so spillovers are unlikely to occur.

The second assumption relates to the absence of anticipation effects and says that a merger is not allowed to affect the market outcomes of the merging firms' brands before the merger. In general, the announcement of a merger does not automatically mean that the merger will be carried out in the future. There are various reasons why a proposed merger may be canceled at a later date. For instance, the merger itself requires negotiations between the owners of the target and the acquirer, and they may fail to reach an agreement. Another possibility is that a due diligence conducted after the merger announcement uncovers problems with the target that the acquirer did not anticipate. Because of all these uncertainties in the period between the merger announcement and the final acquisition, it is unlikely that the retailers will start offering better deals to the targets and/or the acquirers before the merger actually takes place. In this context, it is worth noting that in the consumer packaged goods retail industry, firms often negotiate annually, with some smaller negotiations occurring during the year (for instance, to coordinate the joint marketing and sales effort; see, for instance, Anderson and Fox, 2019 on the planning of trade promotions). This means that the retailers make commitments for a relatively long period and may be less willing to respond to rumors in negotiations.

Another line of reasoning is to think of a counterfactual world in which there were anticipation effects. In this case, anticipation effects would probably be relevant for at most one or two years. Figure 2 shows coefficients covering up to 5 years before the merger, which means that I would expect to see pre-trends. However, since my analysis does not provide any evidence for pre-trends, it renders anticipation effects unlikely.

The third assumption is the existence of parallel trends. Depending on the event study approach used, this assumption comes in different forms, but the general idea underlying this assumption is usually similar. Measuring the average treatment effect on the treated requires comparing the outcome of a treated individual to its untreated counterfactual. This poses a problem since it is obviously not possible to observe an individual in both states (that is, being treated and being untreated) at the same time. Therefore, the empirical strategy is to find an appropriate counterfactual scenario. A simple before-after comparison (that is, a change within an individual over time) would not be suitable since other variables than the treatment status may change, and these variables, when not appropriately controlled for, can introduce a bias. Therefore, the literature usually exploits the presence of untreated or not-yet-treated individuals, with the idea that their outcomes follow a similar development except for the impact of the treatment.

While the idea to compare treated to untreated individuals was originally developed for treatments that occur for all treated individuals at the same time, the literature has also applied this approach to so-called staggered adoptions where the individuals are treated at different points in time. The recent event study literature (in particular Goodman-Bacon, 2021) shows that this has previously unexpected consequences in the sense that researchers may compare groups of individuals to each other that they did not intend to compare. More specifically, treated individuals are compared to other individuals who have been treated earlier. These comparisons are often referred to as forbidden comparisons. These comparisons can lead to a bias if the treatment effects are heterogeneous across cohorts, with a cohort being all firms that are treated in a particular year. This bias can even be strong enough to turn around the sign of an estimate for the average treatment effect on the treated and thus causes serious concerns. In the context of my analysis, the assumption that treatment effects are homogeneous across cohorts is difficult to maintain. In particular, my sample period from 2006 to 2019 covers the financial crisis and the subsequent recovery phase, so mergers of different cohorts also experienced different macroeconomic environments.

Figure 3: Changes in Sales, Quantities and Prices

(a) Sales

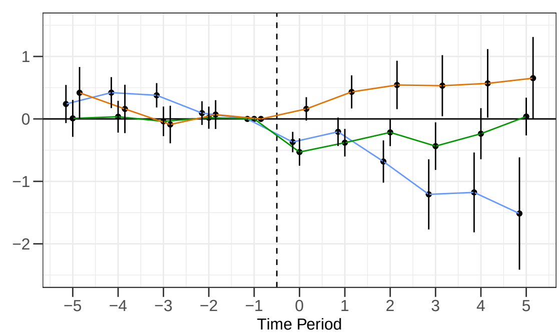

(b) Quantities

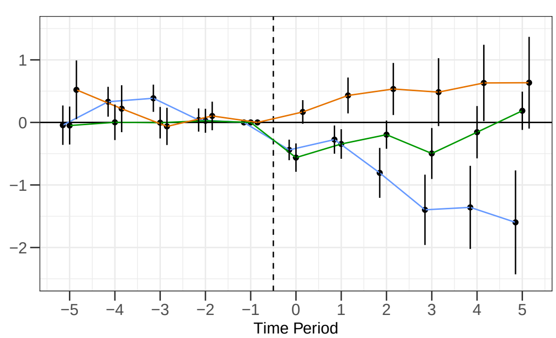

(c) Prices

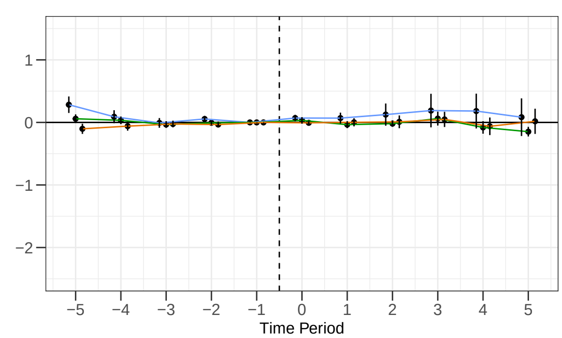

The above considerations are typically discussed in the context of binary treatments. Callaway et al. (2021) highlight that continuous treatments—like in my analysis—further complicate the analysis of causal effects. For instance, they require additional assumptions and stricter versions of some of the previously mentioned assumptions. To simplify my analysis, I convert my continuous measure into a categorical variable. More precisely, I use the logarithms of the pre-merger revenue ratios at the different retailers (as depicted in Figure 1) to split my sample into three groups. I use the 33^{rd} and 66^{th} percentiles (-0.01 and 2.22, respectively) as boundaries. I then apply the estimator of Callaway and Sant'Anna (2021) to each category, treating the treatment variable as binary and ignoring potential variation in the treatment intensity in each subsample. The idea of the approach by Callaway and Sant'Anna (2021) is to perform separate estimations of the treatment effects for each cohort. In each year, the firms that are treated in that year are compared only to those who have not yet been treated and will not receive treatment in the time period used for the comparison. The average treatment effect on the treated is then calculated by a weighted average. Apart from the fact that I use a binary instead of a continuous treatment, another caveat is that the number of mergers, which I previously used as clusters for the standard errors, is already quite small in general, and the number in each subsample is even smaller so that I cluster the standard errors only at the brand level.

Figure 3 shows the result of this exercise. The three lines refer to the different subsamples, with the blue line referring to ratios below the first tertile, the green line to ratios between the first and second tertile, and the orange line to ratios above the second tertile. Panels (a) and (b) show the results for revenues and quantities. Although some pre-merger coefficients are statistically different from 0 at the 5% level, the figures do not show any meaningful pre-trends in general. In contrast, the post-merger coefficients show clear trends that fit the results of my previous analysis. Since the first tertile is roughly 0, the blue line refers to almost all brands for which the targets have larger revenues than the acquirers at a retailer. For these brands, the effect size is negative, and the decrease in revenues 5 years after the merger is roughly -78%. In contrast, targets whose pre-merger revenues are strongly smaller than those of the acquirers at a retailer experience a strongly positive effect (orange line). While the effect is positive in all years, the effect is statistically different from 0 at the 5% level in only some years. Evaluated at the point estimate, the revenues increase by about 92% 5 years after the merger. Finally, if the pre-merger revenues of target and acquirer do not diverge too strongly (green line), there seems to be a slightly negative effect immediately after the merger, but no effect (neither positive nor negative) is visible 5 years after the merger. With respect to changes in prices (Panel (c)), the coefficients are mostly not statistically different from 0 at the 5% level, and even those that are statistically significant are economically negligible in magnitude. Overall, Figure 3 shows similar patterns to those in Figure 2, providing evidence that the previously documented patterns are indeed causal.

Note that Figure 3 also motivates the application of the logarithm to the pre-merger revenue ratios. This is because the direction of the effect takes different signs depending on whether the ratios are above or below 1.

So far, my investigation of the effects of cross-category mergers has focused on the targets. At the end of this section, I will briefly document the effects on the acquirers. To this end,

Figure 4: Changes in Sales, Quantities and Prices (Acquirers)

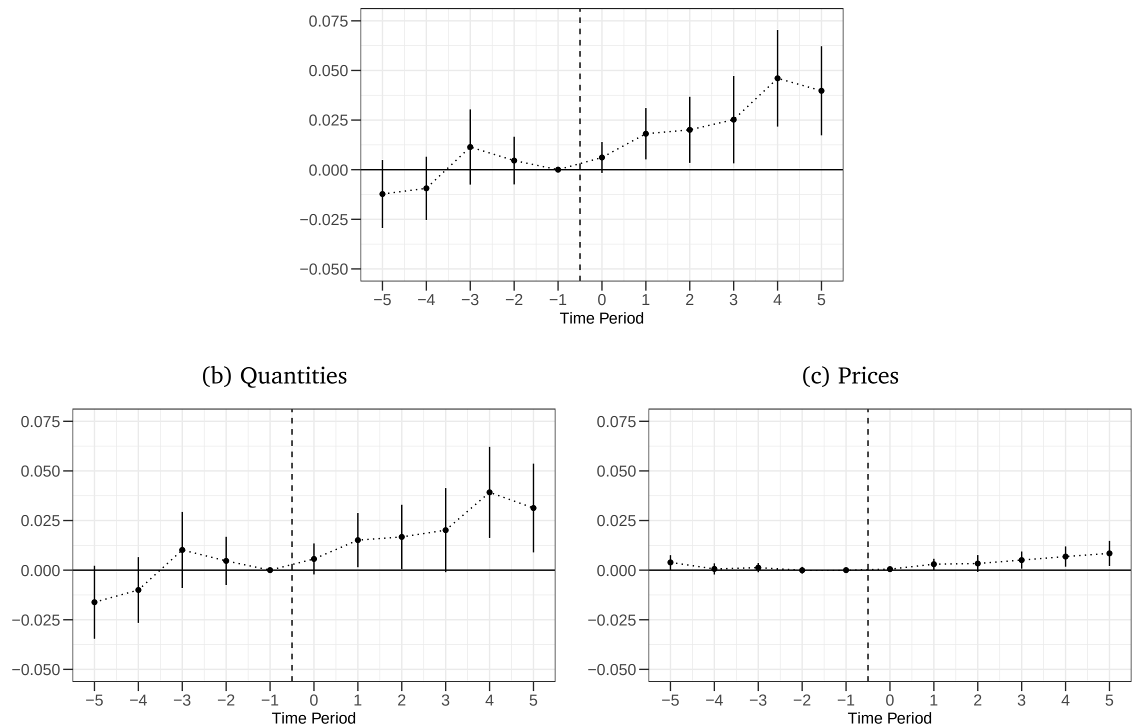

(a) Sales

(b) Quantities

(c) Prices

I estimate a modified version of my baseline specification 2. Since the underlying data now refers to the acquirers' brands, it seems reasonable to adjust the ratio measure and to use the logarithm of the *inverse* ratio. That is, I consider the ratio of the targets' pre-merger revenues to those of the acquirers.

$$\log\left(\frac{total\ sales_{f_j(-1),c,-1}}{total\ sales_{f_j(0),c,-1}}\right) \tag{3}$$

Based on the analysis of the targets, the initial hypothesis is that an acquirer's revenues and quantities increase after the merger if the pre-merger revenues of the acquirer at a retailer are smaller than that of the target, that is if Expression 3 is positive.

Figure 4 is the analog to Figure 2 and shows the results. The bottom line is that the effects go in a similar direction but are of much smaller magnitude. Revenues and quantities increase after the merger, and most of the coefficients are significantly different from 0 at the 5% level. Interestingly, and in contrast to the analysis of the targets, the magnitude of the change in quantities is smaller and thus does not (approximately) match that of the change in revenues (especially in later years). This is because there also seems to be some price effect, although the magnitudes of the corresponding coefficients are still quite small and can be considered economically negligible.

4 Merger Effects on Marginal Costs and Perceived Quality

4.1 Models of Demand and Supply

So far, my analysis has only been concerned with the effect of cross-category mergers on directly observable market outcomes. In the next step, I use the work of Döpper et al. (2023) to shed some light on other measures that can be inferred from the data based on assumptions about demand and firm conduct. To this end, I will first briefly introduce their models of demand and supply and their empirical strategy. In the following, I adopt the notation and formulas from Döpper et al. (2023).

The demand side builds on the seminal work of Berry et al. (1995) and is a random coefficient Logit model. As before, let c denote the retail chain and r the region. The variable t denotes the quarter. A geographic market is defined as a region-retailer combination, which means that the approach abstracts from retailer competition.[20] Since the model will be estimated separately for each category and year, the combined index crt denotes the market level. In each market, consumers can choose between $0, \dots, J_{crt}$ products, where 0 denotes the outside option of not buying any of the products offered by the manufacturers.

Each consumer i is endowed with certain characteristics (like household demographics) and receives an (indirect) utility of u_{ijcrt} when buying product j. The utility of the outside option ($j = 0$) is normalized to zero. The utility of the other products is given by

$$u_{ijcrt} = \beta_i^* + \alpha_i^* \cdot p_{jcrt} + \xi_{jr} + \xi_{cr} + \xi_t + \Delta\xi_{jcrt} + \varepsilon_{ijcrt}, \tag{4}$$

where β_i^* is a consumer-specific constant, p_{jcrt} is the price of product j, α_i^* is a consumer-specific scalar, and ξ_{jr}, ξ_{cr}, and ξ_t are product-region, retailer-region, and quarter fixed effects. $\Delta\xi_{jcrt}$ and ε_{ijcrt} are error terms. The first term is typically called the "structural error term" and captures the reaction of the consumers to unobserved product characteristics, while the second term describes a random consumer-specific taste shock ("Logit" shock). The presence of the structural error term leads to an endogeneity problem when taking the model to the data since unobservable product characteristics might be correlated with observable characteristics like prices.

A key feature of the random coefficient demand model is that it allows for heterogeneity across consumers. Döpper et al. (2023) allow consumers to differ in three characteristics: an unobserved demographic which is standard normal distributed, (the logarithm of) the household income, and a variable indicating whether a household has children ($= 1$) or not ($= 0$).

[20] Note that this does not mean that this approach completely rules out competition between retailers. Instead, the other retailers in the same region are part of the outside option.

The consumer's utility depends on these characteristics through the consumer-specific parameters α_i^* and β_i^*. The consumer-specific constant β_i is allowed to vary in all three characteristics, while the consumer-specific reaction to prices α_i^* is only allowed to vary by the two observable characteristics. Formally, this can be expressed by

$$\alpha_i^* = \alpha + \Sigma_1 D_i \qquad \text{and} \qquad \beta_i^* = \beta + \Sigma_2 D_i + \sigma v_i, \tag{5}$$

where α and β are the mean parameters that are constant across consumers, D_i describes the observable household demographics, v_i is the unobserved demographic, and Σ_1, Σ_2, and σ describe the impact of these demographics (that is, the size of the consumer-specific deviations from the mean parameters).

Each consumer buys the product which yields the highest utility. Taking the random taste shock into account, the choice probability of consumer i buying product j is given by the typical Logit expression

$$s_{ijcrt} = \frac{\exp\left(u_{ijcrt}\right)}{\sum\limits_{k \in 0,\ldots,J_{crt}} \exp\left(u_{ikcrt}\right)}. \tag{6}$$

By integrating over the distributions of consumer characteristics, I can derive the market share s_{jcrt} of product j in market crt. Finally, multiplying the market share with the market size M_{crt}[21] leads to the quantity q_{jcrt} sold of product j in market crt.

The demand model is combined with a supply side. Manufacturers are assumed to set prices to maximize (static) profits, i.e., they compete in static Bertrand competition with differentiated goods. Retailers use a cost-plus pricing strategy and place a constant markup on the prices of the manufacturer. Under this assumption, the retail markup becomes part of the manufacturers' marginal costs, and thus, the approach isolates the manufacturers' markups.

The first-order conditions of the manufacturers' maximization problem lead to the following decomposition of the price:

$$p_{crt} = c_{crt} - \left(\Omega_{crt} \circ \left[\frac{\partial s_{crt}\left(p_{crt}\right)}{\partial p_{crt}}\right]'\right)^{-1} s_{crt}\left(p_{crt}\right), \tag{7}$$

where p_{crt}, s_{crt}, and c_{crt} are vectors capturing prices, market shares, and marginal costs. Ω_{crt} is the ownership matrix with entries in $\{0, 1\}$. If products j and k are owned by the same firm, the entries $[j, k]$ and $[k, j]$ take value 1, otherwise 0. The derivative in square brackets is a matrix that contains the derivative of each market share with respect to each price and thus provides information about the substitution patterns. Finally, $\circ$ denotes the element-wise matrix multiplication (Hadamard product).

[21] Defining the market size across a large number of product categories and years is a non-trivial challenge. See Döpper et al. (2023) for details.

Equation (7) decomposes the price into the marginal cost and the markup. The markup depends on observable variables (like market shares and ownership) and the substitution patterns that can be inferred from the demand side. Therefore, it is possible to calculate the markup for a given set of demand side parameters. Since prices are observed as well, Equation (7) can also be used to calculate the marginal costs directly.

Finally, as noted earlier, $\Delta\xi_{jcrt}$ can include unobserved product characteristics, which can lead to endogeneity problems. The identification strategy will require splitting the marginal costs into an observed and an unobserved component. To this end, the marginal costs are decomposed using product-region (η_{jr}), retailer-region (η_{cr}), and quarter (η_t) fixed effects. The remaining part ($\Delta\eta_{jcrt}$) denotes the unobserved cost shock.

$$c_{jcrt} = \eta_{jr} + \eta_{cr} + \eta_t + \Delta\eta_{jcrt} \tag{8}$$

4.2 Estimation and Identification

The aim is to estimate the unknown parameters of the demand model, which can be divided into two sets. The first set, denoted by Θ_1, contains the parameters α and β.[22] These parameters describe the mean values of α_i^* and β_i^* and can be used to calculate the so-called mean utility δ_{jcrt} by setting the remaining parameters that determine the consumer-specific deviations to zero. In contrast, the second set, denoted by Θ_2, contains all parameters that determine the impact of the consumer characteristics, that is, Σ_1, Σ_2, and σ. For each consumer i and product j, the difference between the utility u_{ijcrt} and the mean utility δ_{jcrt} describes the consumer-specific deviation in the utility space.

Döpper et al. (2023) use a modified version of the nested fixed point estimator of Berry et al. (1995) to estimate both sets of unknown parameters[23]. To understand these modifications, it is useful to consider the mechanics of the estimator first. Consider a given set of candidate parameters for Θ_1 and Θ_2. In the first step, the estimator derives the total utility levels for each product in each market and splits them into the mean utilities and the consumer-specific deviations. To do this, it requires only information about the market shares, the consumer characteristics, and the candidate parameters for Θ_2 (but not the candidate parameters for Θ_1). With these estimates in hand, it is then possible to use the candidate parameters from Θ_1 to further split the mean utility into its components and derive an estimate for the structural error term $\Delta\xi_{jcrt}$. This structural error term is then usually interacted with instrument variables. Most importantly, this procedure shows that the estimator uses the candidate parameters for Θ_1 and Θ_2 in two steps, where each step uses only one set (either Θ_1 or Θ_2).

[22] Technically, ξ_{jr}, ξ_{cr}, and ξ_t are also part of Θ_1. However, the fixed effects are typically not part of the parameters estimated with GMM, but are estimated with Ordinary Least Squares for given candidates for the other parameters.

[23] Apart from the modifications outlined in the following, they adopt some improvements and best practices from Brunner et al. (2017) and Conlon and Gortmaker (2020)

Döpper et al. (2023) use this two-step structure to modify the estimation routine in the following way: In the first step, they use the micro-moments discussed in Section 2 to identify the parameters for Θ_2 that determine the consumer-specific deviations from the mean utility. The idea is that when the model is evaluated at the true parameters, it should predict these micro-moments (see Petrin, 2002; see also Berry and Haile, 2020 and Conlon and Gortmaker, 2023 for further details). Recall that a micro-moment describes the characteristics (like the income) of the average consumer buying a certain product. This means that if consumers behave differently because of their characteristics, this shows up in the micro-moments. For instance, if consumers with low incomes are very price sensitive, the average income of a consumer buying an expensive product should be high. If, in contrast, income does not affect price sensitivity, the average income of a consumer buying an expensive brand should be similar to the average income of the entire population.

For each set of candidate parameters for Θ_2, the micro-moments can be calculated for each product and market, and these predicted moments are then compared to the ones observed in the data. Note that the candidate parameters for Θ_1 are irrelevant for this exercise. To see this, consider the choice probability of a consumer i with certain characteristics buying product j. According to Equation (6), the choice probability depends on the different utility levels that a consumer can achieve when buying the different products (including the outside option). The estimation routine of Berry et al. (1995) can derive these utility levels and split them into mean utilities and consumer-specific deviations. However, as described above, only the parameters from the second set are required to achieve this. The remaining parameters for Θ_1 can split the mean utility into its components, but the total mean utility remains unaffected and does not change in these parameters. If the mean utility is unaffected, the choice probabilities also remain unaffected.

To summarize, Döpper et al. (2023) can use the first step of the estimation routine of Berry et al. (1995) to get an estimate for Θ_2. In the next step, they are concerned with the estimation of the remaining parameters for Θ_1. Since they already have an estimate for Θ_2, they can fix these parameters in the subsequent estimation routine. In particular, this means that the mean utilities are independent of the candidate parameters for Θ_1 and, thus, remain the same.

With the estimates for Θ_2 in hand, Döpper et al. (2023) can derive two measures of interest for given candidate parameters for Θ_1. First, it is straightforward to derive the structural demand-side error term $\Delta\xi_{jcrt}$. To do this, they simply have to subtract the candidate parameter for the constant β and the price multiplied by the candidate parameter for α from the mean utility and then take the fixed effects. Second, with the choice probabilities and the candidate parameter for α, they can calculate the substitution patterns ($\partial s_{crt}\left(p_{crt}\right)/\partial p_{crt}$) required to estimate the marginal costs based on (7). By taking fixed effects, they can then derive an estimate for the cost shock $\Delta\eta_{jcrt}$.

Their key identifying assumption is that the covariance between the two error terms is zero for the true parameter in Θ_1:

$$\operatorname{cov}(\Delta\xi_{jcrt}, \Delta\eta_{jcrt}) = 0. \tag{9}$$

MacKay and Miller (2023) and Döpper et al. (2023) discuss (and justify) this assumption in detail. However, it seems worth briefly pointing out two properties of this approach that make it desirable for application across so many product categories. First, and in contrast to the instrumental variables approach used elsewhere in the literature, covariance restrictions do not require to estimate a first stage. This means that the entire potentially endogenous variation is exploited while instrumental variables restrict the variation. Second, the choice of appropriate fixed effects (or other covariates) is important when deriving estimates for the error terms. The aim is to choose them in a way such that the variation that remains in the error term is unique to this error term. In fact, Döpper et al. (2023) point out that with their fixed effects, the variation that remains in the error term is similar to the one used as instruments elsewhere in the literature. They also achieve similar results when using even stricter fixed effects defined at the product-retailer-region level.

4.3 Merger Effects on Inferred Measures

Döpper et al. (2023) use the empirical strategy outlined in the previous section to estimate the parameters in Θ_1 and Θ_2 separately for all product categories and years. I use their estimates to infer two new measures that cannot be directly observed from the data. First, I calculate the estimated marginal costs $\widehat{c}_{jcrt}$ based on Expression (7). Second, I construct a measure for the perceived quality. To this end, I focus on the mean utility δ_{jcrt} that a consumer gains from buying brand j at retail chain c in region r and quarter t, i.e.,

$$\widehat{\delta}_{jcrt} = \widehat{\beta} + \widehat{\alpha} \cdot p_{jcrt} + \widehat{\xi}_{jr} + \widehat{\xi}_{cr} + \widehat{\xi}_{t} + \Delta\xi_{jcrt}, \tag{10}$$

where the hats indicate estimates. The mean utility has the benefit that it is constant for all consumers and is independent of consumer characteristics. In other words, it omits variables related to horizontal product differentiation so that the remaining utility relates to the price and vertical product differentiation. By subtracting the impact of the price ($\widehat{\alpha} \cdot p_{jcrt}$), I can thus infer a measure of the perceived quality. I regress this measure on brand-retailer-region fixed effects and use the estimated fixed effects in the following. This allows me to exclude seasonal effects because I get an average value at the annual level. Finally, with both new measures in hand, I aggregate the observations across regions and quarters so that I end up with one observation for each brand-retailer-year combination.

It seems reasonable to briefly provide some intuition for the measure of perceived quality. In particular, it is important to highlight that the perceived quality of a brand may differ from the actual quality. As indicated by the name, it depends on the consumers' opinion of the product. This opinion may change, for instance, if a brand is heavily advertised. Another factor could be the shelf space that a retailer allocates to a brand. If the brand occupies a lot of shelf space, this might be a sign that the retailer "believes" in the high potential of a brand. These (mostly psychological) factors are not explicitly modeled in the demand model, and identifying each of them is potentially challenging on its own. Thus, I consider the fixed effects that enter the mean utility, and hence the perceived quality, as a "reduced-form" approach to capturing the average consumer's opinion of a brand while remaining silent about the psychological channels that lead to such an opinion.

I re-estimate the baseline specification 2 using the new measures as dependent variables. Figure 5 shows the result of this exercise. The panels in the first row refer to the targets, and those in the second row refer to the acquirers. In line with the previous analysis, I also use the adjusted ratio measure (3) for the acquirers.

Figure 5: Changes in Marginal Costs and Perceived Quality

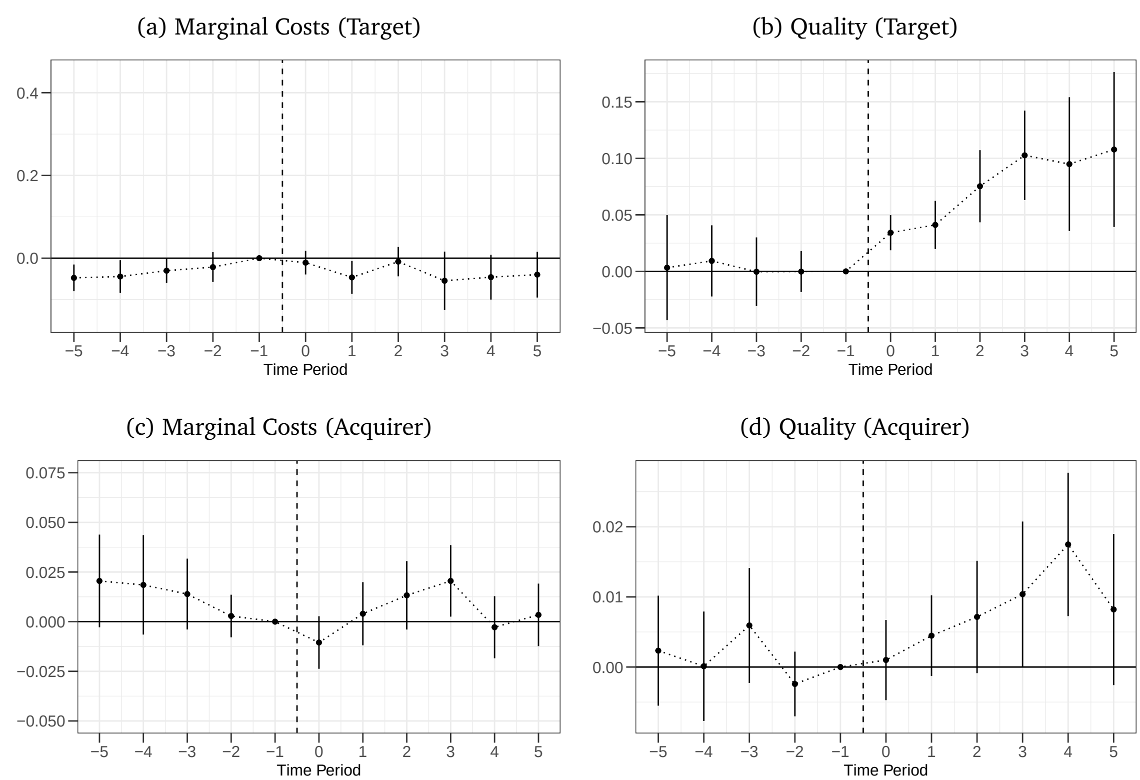

(a) Marginal Costs (Target)
(b) Quality (Target)
(c) Marginal Costs (Acquirer)
(d) Quality (Acquirer)

The first column of Figure 2 shows the results for the marginal costs. Note that I use the same scaling for the y-axis as in Figure 2 and 4, respectively, to simplify the comparison across the graphs for the reader. Panel (a) shows that the marginal costs of the targets are, by and large, unaffected by the merger. Although some coefficients are statistically different from 0 at the 5% level, the magnitudes are so small that they can be considered economically negligible. Interestingly, this is different for the acquirers. The coefficients for changes in the marginal costs of the acquirers show a larger dispersion, but almost all coefficients (except the one in $\tau = 3$) are statistically insignificant at the 5% level. While there seems to be no clear post-merger trend, the pre-merger marginal costs show (if any) a falling trend over time.

The second column shows the results for the perceived quality. Since quality is defined as the sum of the fixed effects, which can take both positive and negative values in general, I use a standardized version of the measure rather than the logarithm as a dependent variable.[24] Therefore, I also do not follow the convention to keep the scaling of the y-axis consistent with those in the other figures.

The results show that perceived quality stays relatively constant before the merger and increases afterward. While this pattern is rather sharp for the targets, meaning that the pre-merger coefficients are extremely close to 0 and that the post-merger coefficients are all statistically different from 0 at the 5% level, the pattern for the acquirers is more noisy. In particular, most post-merger confidence intervals contain 0, even though it is usually close to the interval boundaries. Overall, however, both panels show trends that match the development of the revenues and quantities.

5 Potential Mechanisms and the Portfolio Power Theory

The previous sections paint a rather clear picture of the impact of cross-category mergers on market outcomes. First, cross-category mergers can influence the revenues of both targets and acquirers. The direction of the effect depends on the relative size of the acquirer's revenues to the target's revenues at a given retailer. Second, changes in revenues are almost exclusively driven by changes in quantities and not by changes in prices. Third, by making use of the work of Döpper et al. (2023), I find that marginal costs are almost unaffected and that the changes in quantities can be rationalized within a structural model by changes in the non-price utility part.

The last section is now devoted to a brief discussion of mechanisms that could potentially drive the results. I start with two mechanisms that can be subsumed under the portfolio power theory. The first channel deals with the manufacturers' bargaining power. The idea is that a firm's bargaining power and its ability to influence bargaining outcomes in its own interest depends on its importance for the other firm's profit. This idea is formalized in the Nash-in-Nash bargaining framework, which is frequently used by economists in empirical studies of bargaining (see Draganska et al., 2010; Noton and Elberg, 2018 for examples from the consumer packed goods retail industry and Collard-Wexler et al., 2019 for a micro-foundation). Think of a manufacturer f and a retailer r conducting negotiations over some form of financial

[24]For the standardized version, I first subtract the mean and then divide the measure by its standard deviation.

payments (for instance, linear prices or fixed fee payments) captured in a vector μ_{fr} and an effort level e_{fr} that describes the effort that a retailer spends on the product of manufacturer f (like shelf space or in-store promotions). Then, according to the Nash-in-Nash bargaining framework, they choose these variables to maximize the following expression:

$$\left(\pi_f\left(\mu, e\right) - \pi_f^{-r}\left(\mu^{-r}, e^{-r}\right)\right)^{\lambda}\left(\pi_r\left(\mu, e\right) - \pi_r^{-f}\left(\mu^{-f}, e^{-f}\right)\right)^{1-\lambda} \tag{11}$$

π_f and π_r refer to the profits of the manufacturer and the retailer, and μ and e are vectors capturing the strategic variables of all bargaining pairs. The parameter $\lambda \in (0,1)$ captures the other determinants influencing the abilities of the parties in the negotiations, like, for instance, the negotiation skills of the managers. The superscripts $-f$ and $-r$ refer to a situation where the bargaining between the firms breaks down so that retailer r does not sell products of manufacturer f.

The brackets in Expression 11 show the so-called gains from trade, that is, the extra profit that a firm gains through a collaboration with the other firm. A firm gains bargaining leverage over the other firm if the extra profit for the other firm increases. Intuitively, if the products of the other firm are very important for one's own revenues, a bargaining breakdown would be very costly, and the incentive to settle the negotiation increases.

A merger between two manufacturers leads to a change in the gains from trade since the merging firms are now negotiating jointly with the retailers. Before the merger, a bargaining breakdown with one manufacturer still allows a retailer to settle the negotiations with the other manufacturer. However, this is not possible after the merger, and a bargaining breakdown will result in a loss of the products of both firms. This gives the integrated manufacturer a larger bargaining leverage. Dafny et al. (2019) use a similar reasoning in their analysis of cross-market hospital mergers.

There might also be other determinants than the gains from trade that play a role in the negotiations and that could be affected by a merger. In particular, a merger can alter the logistics of the firms and allow them to operate a better distribution network. One key consideration could be the ability of manufacturers to reliably manage deliveries. For instance, if the product of a rather small manufacturer is subject to highly volatile demand and the manufacturer cannot operate a large distribution network due to its size, the retailer might run out of the product and the shelf space remains empty whenever a demand spike occurs. An alternative strategy would be to increase the inventory, leading to increased costs for the retailer. This gives the retailer small incentives to provide the manufacturer with more shelf space.[25] A merger could give the manufacturers the ability to combine their distribution networks. Apart from potential (fixed) cost savings, this might increase both the reliability of regular deliveries (e.g., due to

[25]The marketing and operations research literature has devoted an entire subfield to the question of optimal shelf space allocation and, hence, forgone profits due to stock-outs have long been an important topic (see Curhan, 1973; Gilliver and Gordon, 1978; Emmelhainz and Stock, 1991 for examples of early studies on this topic).

more frequent deliveries and larger truck loads) and the ability to react to irregular delivery requests. These improvements likely depend on manufacturers' past relationships with a retailer since the logistical operations have likely developed to serve retailers that were willing to sell many products from the manufacturers in the past.

Both channels, the bargaining power channel and the improvements in the distribution network, fit the previously documented patterns in that they depend on the manufacturer-retailer-specific relationships. Both also fit the portfolio power theory because they do not depend on the substitutability/complementarity of the products but on the size of the total sales to a retailer. There are two other explanations that are not related to the portfolio power theory and that I deem less likely to explain the patterns. Both explanations have in common that a merger can lead to better access to resources, in particular financial and human resources.

The first explanation is that a merger leads to increased marketing expenditures. As discussed earlier, marketing activities can also be part of the negotiations between manufacturers and retailers (captured by the effort variable in (11)). Hence, I relate here to retailer-independent marketing activities. Such activities can change the perception of consumers about the quality of the products of the merging parties (for instance, due to stronger brand preferences). Although this channel might seem to fit the patterns at first glance, there is a good reason to believe that this is actually not the case. This is because my analysis aims to isolate retailer-specific effects and that retailer-independent marketing activities would have affected all retailers in a similar manner.[26]

The same argument applies to potential efficiency gains beyond the previously described improvements in the distribution network. Such efficiency gains can include improvements in the production process due to knowledge spillovers or better access to financial resources that spur investments in new technologies. Efficiency gains are part of a longstanding discussion on the competitive and anti-competitive effects of (horizontal and vertical) mergers, dating back to at least Williamson (1968) (see Affeldt et al., 2021 and the references therein for a recent overview).

In the context of my investigation, there is no evidence for investments in (marginal) cost-reducing production technologies since marginal costs do not fall after the merger. This does, however, not necessarily mean that efficiency gains are absent. Another explanation could be that improvements in production technology lead to quality upgrades at the same or similar marginal cost levels. However, I can apply the same argument that renders effects through increased retailer-independent advertising spending unlikely. If the quality improves, the non-price part of consumers' utility will increase, but I would expect this increase to be rather similar across retailers, independent of the target's or acquirer's historical revenues to a retailer.

To summarize, out of the four possible mechanisms discussed, only two mechanisms seem

[26]As mentioned in Section 3.3, the estimated $\beta_{1\ell}$ coefficients from my main specification 2, which are intended to capture retailer-independent effects, are always close to 0 and negligible in magnitude.

to be able to explain the pattern described in my analysis. These two channels can also be subsumed under the portfolio power theory.

6 Conclusion

This paper documents the presence, direction, and size of portfolio effects by analyzing 57 consummated mergers of manufacturers in the US consumer packaged goods retail industry between 2006 and 2019. My analysis focuses on cross-category mergers where the merging firms have (almost) no overlap in their product portfolio prior to the merger. I exploit the large heterogeneity in the pre-merger bargaining positions of the targets and the acquirers at the different retailers (as measured by their pre-merger revenues at the respective retailers) and provide evidence that manufacturers with weaker pre-merger bargaining positions benefit from cross-category mergers through increases in revenues, while manufacturers with stronger pre-merger bargaining positions are harmed and experience revenue decreases. These increases (decreases) in revenues are almost entirely driven by increases (decreases) in the quantities sold and not by changes in prices. I show that these patterns can be rationalized within a structural model by changes in the perceived quality of the products of the merging firms. Changes in marginal costs do not seem to play a crucial role.

In the last section, I discussed two potential mechanisms related to the portfolio power theory that help explain these patterns. Both build on the idea that an increase in the sheer size of the product portfolio can impact the negotiations between the manufacturers and the retailers. The first channel is that bargaining breakdowns become increasingly costly for a retailer when the size of the manufacturer increases. This changes the incentives of the retailers to settle the negotiations with the manufacturers and allows the manufacturers to demand larger concessions (for instance, in the form of better or more shelf space). The second channel builds on improvements in logistics because the merging firms can operate a joint distribution network. The better logistics increase the incentives for the retailers to provide the products of the merging firms with more and better shelf space since stockouts (or similar problems) are less likely to occur. Finally, I argue that changes in advertising spending and efficiency gains are unlikely to explain the patterns.

An open question that I cannot answer at the moment is that of possible policy implications. To shed light on this question, I plan to use a structural model in future versions of this paper that will serve two purposes. On the one hand, it provides me with insights into the impact of portfolio effects on consumer surplus and welfare; on the other hand, it allows me to investigate whether portfolio effects can be regarded as pro- or anti-competitive in the merger cases I study.

Appendices

A Estimates of the Two-Way Fixed Effects Regressions

	log(Sales)	log(Market Share)	log(Quantity)	log(Price)	log(Marginal Cost)	Sd. Quality
Ratio in t = -5	0.012	0.009	0.022	-0.011***	-0.048**	0.003
	(0.039)	(0.041)	(0.039)	(0.003)	(0.016)	(0.023)
Ratio in t = -4	-0.005	-0.009	0.001	-0.005	-0.044*	0.009
	(0.034)	(0.035)	(0.035)	(0.003)	(0.020)	(0.016)
Ratio in t = -3	-0.015	-0.018	-0.013	-0.003	-0.030*	0.000
	(0.033)	(0.033)	(0.033)	(0.002)	(0.015)	(0.015)
Ratio in t = -2	0.010	0.011	0.011	-0.001	-0.022	0.000
	(0.022)	(0.022)	(0.022)	(0.001)	(0.018)	(0.009)
Ratio in t = 0	0.086***	0.083***	0.086***	0.000	-0.011	0.034***
	(0.015)	(0.015)	(0.014)	(0.001)	(0.014)	(0.008)
Ratio in t = 1	0.112***	0.114***	0.114***	-0.002	-0.046*	0.041***
	(0.024)	(0.025)	(0.025)	(0.003)	(0.020)	(0.011)
Ratio in t = 2	0.187***	0.182***	0.188***	-0.001	-0.008	0.075***
	(0.036)	(0.038)	(0.038)	(0.004)	(0.018)	(0.016)
Ratio in t = 3	0.235***	0.233***	0.239***	-0.004	-0.055	0.103***
	(0.048)	(0.050)	(0.049)	(0.005)	(0.035)	(0.020)
Ratio in t = 4	0.260***	0.267***	0.272***	-0.012*	-0.046	0.095**
	(0.066)	(0.070)	(0.066)	(0.005)	(0.027)	(0.030)
Ratio in t = 5	0.306***	0.309***	0.317***	-0.011*	-0.040	0.108**
	(0.074)	(0.079)	(0.073)	(0.005)	(0.028)	(0.034)
Ratio in t<-5	0.007	0.004	0.020	-0.013*	-0.058**	-0.007
	(0.051)	(0.053)	(0.050)	(0.006)	(0.022)	(0.024)
Ratio in t>5	0.288***	0.265***	0.301***	-0.014**	-0.003	0.097**
	(0.061)	(0.060)	(0.060)	(0.005)	(0.020)	(0.029)
Num.Obs.	62537	62537	62537	62537	53615	62468
R2	0.757	0.686	0.818	0.986	0.876	0.656
R2 Adj.	0.728	0.649	0.797	0.985	0.859	0.616
Retailer-Merger FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
Period FE	X	X	X	X	X	X

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B Effects on Revenue and Quantity Shares

Figure 6: Changes in Revenue and Quantity Shares

(a) Sales Shares

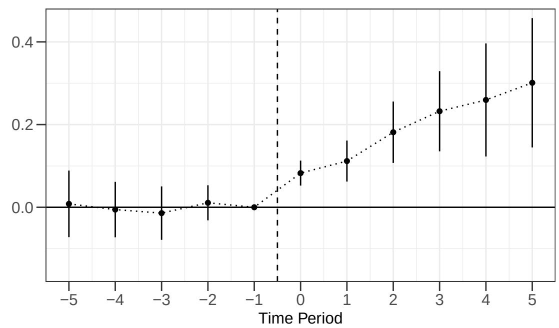

(b) Quantity Shares

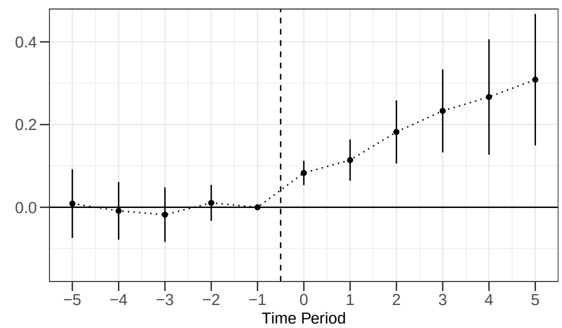

C Results with Balanced Panel

Figure 7: Changes in Revenues, Quantities, and Prices (Balanced Panel)

(a) Sales

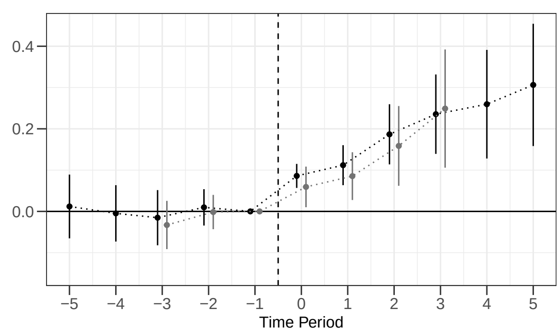

(b) Quantities

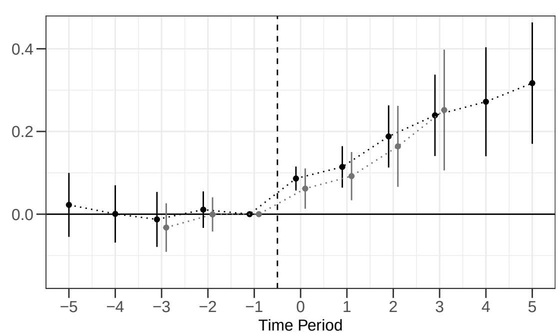

(c) Prices

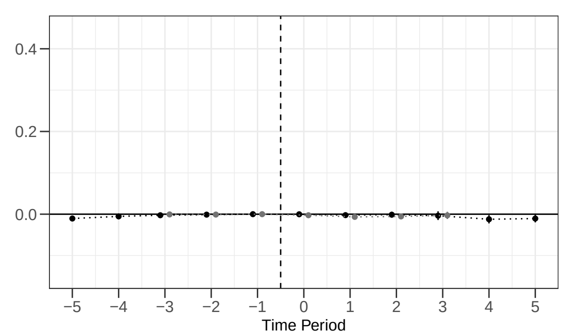

D Results without Time Period Fixed Effects

Figure 8: Changes in Revenues, Quantities, and Prices (without Time Period Fixed Effects)

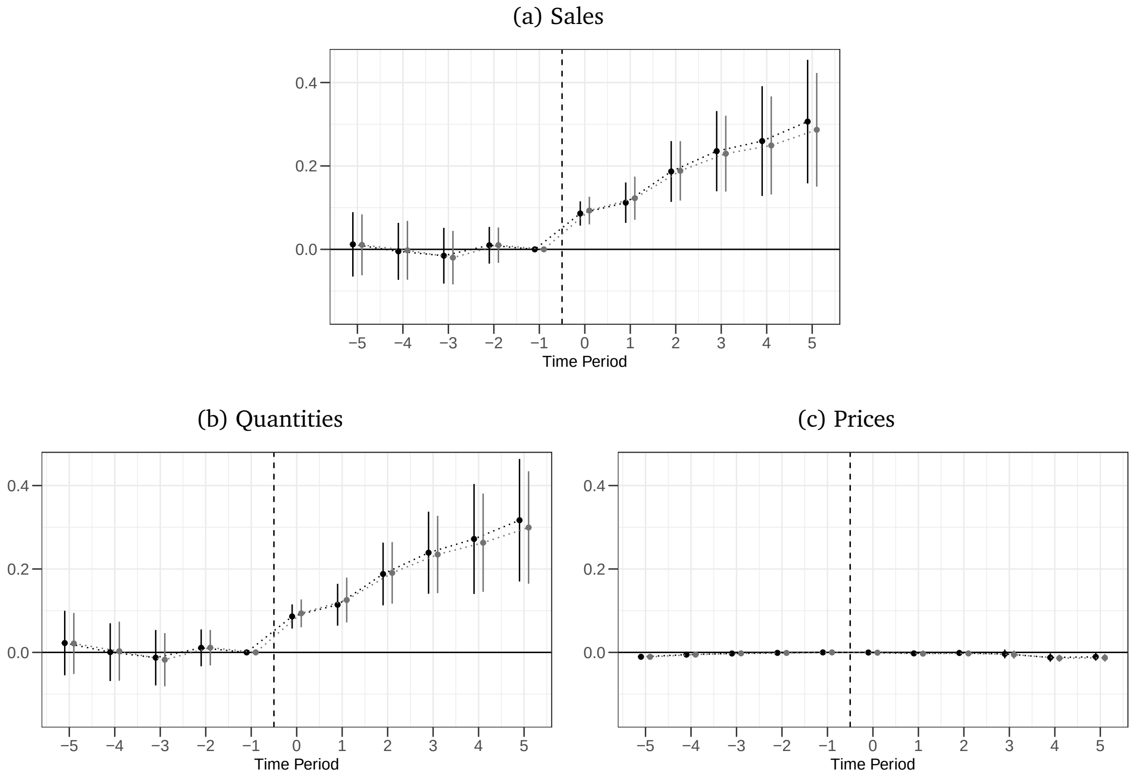

References

AFFELDT, P., T. DUSO, K. GUGLER, AND J. PIECHUCKA (2021): "Assessing EU Merger Control through Compensating Efficiencies," *DIW Berlin: Discussion Papers 1979*.

ANDERSON, E. T. AND E. J. FOX (2019): "How price promotions work: A review of practice and theory," *Handbook of the Economics of Marketing*, 497–552.

ASHENFELTER, O. AND D. HOSKEN (2010): "The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin," *The Journal of Law & Economics*, 53.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

BERRY, S. T. AND P. A. HAILE (2020): "Nonparametric Identification of Differentiated Products Demand Using Micro Data," *NBER Working Paper 27704*.

BHATTACHARYA, V., G. ILLANES, AND D. STILLERMAN (2023): "Merger Effects and Antitrust Enforcement: Evidence from US Retail," *NBER Working Paper 31123*.

BRUNNER, D., F. HEISS, A. ROMAHN, AND C. WEISER (2017): "Reliable Estimation of Random Coefficient Logit Demand Models," *DICE Discussion Paper No. 267*.

CALLAWAY, B., A. GOODMAN-BACON, AND P. H. SANT'ANNA (2021): "Difference-in-Differences with a Continuous Treatment," *Mimeo (available via arXiv:2107.02637)*.

CALLAWAY, B. AND P. H. SANT'ANNA (2021): "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 225, 200–230.

CHOI, J. P. (2001): "A Theory of Mixed Bundling Applied to the GE/Honeywell Merger," *Antitrust*, 16, 32–33.

CHUNGA, J. AND S. JEON (2014): "Portfolio effects in conglomerate mergers: the empirical evidence of leverage effects in Korean liquor market," *Applied Economics*, 46, 4345–4359.

COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. S. LEE (2019): ""Nash-in-Nash" Bargaining: A Microfoundation for Applied Work," *Journal of Political Economy*, 127, 163–195.

CONLON, C. AND J. GORTMAKER (2020): "Best practices for differentiated products demand estimation with PyBLP," *RAND Journal of Economics*, 51, 1108–1161.

——— (2023): "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP," *Mimeo*.

CURHAN, R. C. (1973): "Shelf Space Allocation and Profit Maximization in Mass Retailing," *Journal of Marketing*, 37, 54–60.

DAFNY, L., K. HO, AND R. S. LEE (2019): “The price effects of cross-market mergers: theory and evidence from the hospital industry,” *RAND Journal of Economics*, 50, 286–325.

DE CHAISEMARTIN, C. AND X. D’HAULTFŒUILLE (2023): “Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey,” *Econometrics Journal*, 26, C1–C30.

DÖPPER, H., A. MACKAY, N. MILLER, AND J. STIEBALE (2023): “Rising Markups and the Role of Consumer Preferences,” *Kilts Center at Chicago Booth Marketing Data Center Paper Series*.

DRAGANSKA, M., D. KLAPPER, AND S. B. VILLAS-BOAS (2010): “A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel,” *Marketing Science*, 29, 57–74.

DUBÉ, J.-P., A. HORTAÇSU, AND J. JOO (2021): “Random-Coefficients Logit Demand Estimation with Zero-Valued Market Shares,” *Marketing Science*, 40, 637–660.

EMMELHAINZ, M. A. AND J. R. STOCK (1991): “Consumer Responses to Retail Stock-outs,” *Journal of Retailing*, 67, 138–147.

EVANS, D. S. AND M. SALINGER (2002): “Competition Thinking at the European Commission: Lessons from the Aborted GE/Honeywell Merger,” *George Mason Law Review*, 10, 489–532.

GANDHI, A., Z. LU, AND X. SHI (2023): “Estimating demand for differentiated products with zeroes in market share data,” *Quantitative Economics*, 14, 381–418.

GILLIVER, A. AND H. A. GORDON (1978): “An Analytic Information System for a Representative Sales Force – A Case Study,” *Journal of the Operational Research Society*, 29, 719–730.

GOODMAN-BACON, A. (2021): “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 225, 254–277.

GUADALUPE, M., O. KUZMINA, AND C. THOMAS (2012): “Innovation and Foreign Ownership,” *American Economic Review*, 102, 3594–3627.

HENDEL, I. AND A. NEVO (2006): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74, 1637–1673.

LEWIS, M. S. AND K. E. PFLUM (2017): “Hospital systems and bargaining power: evidence from out-of-market acquisitions,” *RAND Journal of Economics*, 48, 579–610.

MACKAY, A. AND N. H. MILLER (2023): “Estimating Models of Supply and Demand: Instruments and Covariance Restrictions,” *Harvard Business School Strategy Unit Work Paper No. 19-051 (first version: 2018, revised version: August 2023)*.

Majerovitz, J. and A. Yu (2023): "Consolidation on Aisle Five: Effects of Mergers in Consumer Packaged Goods," *Mimeo (available via SSRN)*.

Noton, C. and A. Elberg (2018): "Are Supermarkets Squeezing Small Suppliers? Evidence from Negotiated Wholesale Prices," *Economic Journal*, 128, 1304–1330.

OECD (2001): "Portfolio Effects in Conglomerate Mergers," *OECD Journal: Competition Law and Policy*, 4.

Park, S. (2009): "An empirical testing of leverage effects via the common distribution network," *International Review of Law and Economics*, 29, 143–152.

Patterson, D. E. and C. Shapiro (2001): "Transatlantic Divergence in GE/Honeywell: Causes and Lessons," *Antitrust*, 16, 18–26.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705–729.

Reynolds, R. J. and J. A. Ordover (2002): "Archimedian Leveraging and the GE/Honeywell Transaction," *Antitrust Law Journal*, 70, 171–198.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 235, 2218–2244.

Watson, P. (2003): "Portfolio effects in EC merger law," *Antitrust Bulletin*, 48, 781–806.

Williamson, O. E. (1968): "Economies as an Antitrust Defense: The Welfare Tradeoffs," *American Economic Review*, 58, 18–36.